

BACKGROUND OF THE INVENTION

1. Cross-Reference to Related Application

5 This application relates to co-pending U.S. patent application, serial no. 09/624,850 (Attorney Reference No: 5181-55000), filed on the same day as the present application and entitled "System and Method for Implementing a Routing Scheme in a Computer Network Using Intention Packets When Fault Conditions Are Detected" by Whay S. Lee, which is incorporated herein by reference in its entirety, and to co-pending
10 U.S. patent application, serial no. 09/624,854 (Attorney Reference No: 5181-63400), filed on the same day as the present application and entitled "Routing Mechanism Using Intention Packets in a Hierarchy of Networks" by Whay S. Lee, which is incorporated herein by reference in its entirety.

15 2. Field of the Invention

This invention relates generally to the field of computer systems and, more particularly to communications protocols within computer systems and/or networks, and communication routing or switching within computer systems and/or networks.

20

3. Description of the Related Art

25 Devices in a communications network may communicate with one another by sending packets to each other. These packets may correspond to an operation that one device seeks to perform with the assistance of another device. Packets may be transferred between devices using a network which may include a fabric. A network generally may comprise a number of devices that are configured to convey packets between any number of devices.

005220 T8E2960

At times, however, adverse transmission conditions may occur in the network when one device is sending a packet to another device on the network. For example, the network may become overly congested with packets traveling between devices. Another
5 adverse condition that may occur in networks is corruption of a packet (e.g. an ECC error, parity error, checksum error, etc.). A fault condition (e.g. failure in the fabric) that prevents a packet from proceeding along its proposed route is another adverse condition. One adverse condition resolution scheme in packet routing protocols involves dropping packets in response to an adverse transmission condition within one or more devices (e.g.
10 a switch, router, endpoint, node, etc.) in the network. In the case of congestion, a device within the network may drop one or more packets randomly or selectively to instantaneously reduce the packet load. By dropping one or more packets, the device can free up sufficient resources to service the remaining packets. Corrupted packets may also be dropped. For a fault condition, all packets reaching a fault point may have to be
15 dropped and resent using a different route.

When a packet is dropped, the device that originated the packet needs to detect that it has been dropped so that it can be resent. Detection of dropped packets may be accomplished by a time-out mechanism at a receiving device when it fails to receive an
20 expected packet, by a time-out mechanism at a sending device where the sending device is waiting for a signal from a receiving device that a packet has been received, or by sending a negative acknowledgement (NACK) from the device that dropped the packet to the sending device.

25 Each of these detection methods has its problems, however. Where a time-out mechanism is used, a device needs to ensure that it does not resend a packet until it is certain that the original packet has been dropped. The device must generally wait a period of time to ensure that the packet has been dropped and this waiting period results

in a performance penalty. Where a NACK to the sending device is used to signal a dropped packet, a dropped packet is unknown to the receiving device, no matter how far through the network the dropped packet traveled before being dropped. The NACK propagates back only to the sender and the network resources consumed by the dropped packet are wasted. An apparatus and method of handling adverse transmission conditions that minimizes the drawbacks of the above methods is needed.

SUMMARY

The problems outlined above may be solved by the use of the apparatus and method described herein. Generally speaking, an apparatus and method for implementing a routing scheme using intention packets is contemplated. At times, a switching device within a network may detect an adverse transmission condition such as a fault, corruption, or congestion. When a switching device detects an adverse transmission condition, it may drop one or more packets to alleviate the adverse condition.

For example, in the case of congestion, the switching devices may support a particular amount of routing resource (e.g. bandwidth) in and out of each of their ports. When a packet or a header portion of a packet arrives at a switching device, the switching device may determine what port the packet will need and the amount of resources required by the packet on that port. If the required resources are available for the packet on the port, then the switching device may route the packet to a next device. If the required resources are not available for the packet on the port, then the switching device may drop at least a portion of the packet. A switching device may be configured to remember the amount of resources that it has committed to other traffic and may drop a subsequent packet when it is unable to guarantee the resources requested by the subsequent packet.

When a packet is dropped due to an adverse transmission condition, the switching device that dropped it may create a second packet to be forward to the destination device. The second packet may be an intention packet that corresponds to the dropped packet. The switching device may then route or forward the intention packet to the destination device specified by the dropped packet. The intention packet may be much smaller than the dropped packet and may include a header portion or part of a header portion of the dropped packet, as discussed below. Accordingly, it may require substantially less

00527072500

Sub
A1

AI Cont.

bandwidth than the dropped packet so that a switching device may be able to guarantee its delivery to the next device if congestion was encounter. For fault conditions, the device may send the intention packet to the destination using a different route in order to avoid the fault condition.

5

A device may react to receiving an intention packet based on the information in the intention packet as well as its own characteristics. The information in the intention packet may indicate whether the dropped packet was a request or a response to a request. If the dropped packet was a request, then the device that receives the intention packet, i.e. the receiving device, may send a signal to the requesting device that its request has been dropped. The receiving device may also perform an operation that corresponds to the request in anticipation of receiving the resent packet. If the dropped packet was a response to a request, then the requesting device will receive the intention packet. The requesting device may reschedule the request in response to receiving the intention packet and may convey a subsequent packet that includes the request to the receiving device. The requesting device may be configured optimize the resending of the request by rescheduling the request for a more appropriate time or by merging the request with a subsequent request.

10

15

20

00527078252922

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram illustrating one embodiment of devices connected through a network.

5

Fig. 2a is a block diagram illustrating one embodiment of a packet.

Fig. 2b is a block diagram illustrating one embodiment of a packet.

10

Fig. 3 is a flow chart illustrating one embodiment of a method for implementing an intention packet routing scheme in a switching device.

Fig. 4 is a flow chart illustrating one embodiment of a method for handling packets at a first exemplary device.

15

Fig. 5 is a flow chart illustrating one embodiment of a method for handling packets at a second exemplary device.

Fig. 6 is a block diagram of one embodiment of a system including one or more networks.

20

Fig. 7 is a block diagram of one embodiment of a plurality of nodes connected in a network.

While the invention is described herein by way of example for several embodiments and illustrative drawings, those skilled in the art will recognize that the invention is not limited to the embodiments or drawings described. It should be understood, that the drawings and detailed description thereto are not intended to limit the

25

invention to the particular form disclosed, but on the contrary, the intention is to cover all modifications, equivalents and alternatives falling within the spirit and scope of the present invention as defined by the appended claims.

5

09625331.072500

DETAILED DESCRIPTION OF AN EMBODIMENT

Turning now to Fig. 1, a block diagram illustrating one embodiment of devices connected through a network. Fig. 1 depicts device 100 coupled to device 120 through network 110. Network 110 includes switching devices 112a through 112(n). It can be noted that any number, n, of switching devices may be included in network 110 and that the switching devices may be connected in any manner. It can be noted that device 100 and device 120 can comprise any type of devices configured to communicate with other devices. In one particular embodiment, device 100 may comprise a computer system host interface to a storage system and device 120 may comprise a disk drive. Alternatively, devices 100 and 120 may be any combination of networked devices, such as client computers, servers, storage devices, etc. Network 110 may be a network on any scale, such as from the Internet, to a LAN or a network inside a single system interconnecting multiple disk drives or other I/O devices. In an Internet embodiment, device 100 may be a requestor of a web page and device 120 may be a web server.

Devices 100 and 120 may be configured to communicate through network 110. Device 100 may convey a request to switching device 112a. The request may be routed or transmitted through network 110 until it is received at switching device 112(n) and ultimately at device 120. Device 120, depending on the request, may convey a response to the request to switching device 112(n). The response may be routed through network 110 until it is received at switching device 112a and ultimately at device 120. In one embodiment, switching devices 112 are simple switches that include at least one input port, one output port, and operate to pass data from an input port to an output port. In other embodiments, switching devices 112 may be more intelligent devices such as packet routers or computer systems.

Generally speaking, a request may be a signal from a first device to a second device for the second device to perform an operation or operations. For example, a request may signal a device to read data from or write data to a particular location. A response may be a signal from the second device to the first device that corresponds to the request. A response, for example, may indicate whether or not the operation signaled by the request was performed successfully or may convey data corresponding to the request. In one embodiment, requests and responses may be conveyed using packets defined by a communications protocol. Devices 100 and 120 may be configured to convert their respective requests and responses into packets prior to conveying the packets to network 110.

The switching devices in network 110 may implement a transmission scheme using intention packets. At times, a switching device 112 within a network 110 may detect an adverse transmission condition such as a fault, corruption, or congestion. When a switching device 112 detects an adverse transmission condition, it may drop one or more packets to alleviate the adverse condition.

Switching devices 112a through 112(n) may each include one or more ports. The switching devices may support a particular bandwidth (and/or other routing resources) in and out of each of their ports. When a packet or a header portion of a packet arrives at a switching device, the switching device may determine what port the packet will need and the amount of routing resources required by the packet on that port. If the required routing resources are available for the packet on the port, e.g. the required resources are less than or equal to the available resources, then the switching device may route the packet to a next device. If the required resources are not available for the packet on the port, then the switching device may drop at least a portion of the packet. A switching device may be configured to remember the amount of routing resources that it has committed to other traffic, i.e. other packets, and may drop a subsequent packet when it is

unable to guarantee the resources requested by the subsequent packet. A switching device may also drop corrupted packets, e.g. packets with a ECC, parity, or checksum error. A switching device may also drop packets for other adverse conditions such as a fault condition that prevents a packet from being routed along its proposed route.

5

When a packet is dropped, the switching device that dropped it may create an intention packet that corresponds to the dropped packet. The switching device may then route the intention packet to the destination device specified by the dropped packet. The intention packet may be much smaller than the dropped packet and may include a header portion or part of a header portion of the dropped packet, as discussed below. Accordingly, it may require substantially less routing resources than the dropped packet so that a switching device may be able to guarantee its delivery to the next device. In some cases, such as fault conditions, alternate paths may have to be tried for the intention packet to reach its destination.

15

Fig. 2a illustrates one embodiment of a packet 200 that may be used in the embodiment of Fig. 1. Packet 200 may include a header 202 and a payload 204. Header 202 may include sequencing and network routing information, indicators of the source and destination devices, an address corresponding to a beginning data block, and an indicator of the number of data blocks involved in the request that corresponds to the packet. Payload 204 may include one or more data blocks being transferred by packet 200. The number of data blocks in payload 204 may be variable or fixed according to a given communications protocol. It can be noted that payload 204 may be empty (or absent) for certain types of packets. Header 202 and/or payload 204 may include other information as required by a protocol.

25

In one embodiment, network 110 may implement a bandwidth-reserving wormhole routing model. In this model, header 202 may travel ahead of its

corresponding, variable-size payload 204 in network 110. As the header routes from switching device to switching device in network 110, it may reserve bandwidth (and/or other resources) for its payload 204 at each subsequent switching device. Once the resource are reserved in a switching device, the switching device can assure that payload 204 will route successfully to the next switching device. If a switching device along the way cannot provide the required resources for payload 204, then that switching device may be configured to drop payload 204 and generate an intention packet as discussed above. In this embodiment, the switching device may generate the intention packet using header 202 or part of header 202.

10

Fig. 2b illustrates one embodiment of an intention packet 222. Intention packet 222 may include intention information 226, severed flag 228 and diagnostic information 230. In one embodiment, intention information 226 may include the information of header 202 of the dropped packet. In other embodiments, intention information 226 may include only selected or modified portions of header 202 of the dropped packet. Severed flag 228 may include an indicator that the intention packet corresponds to a dropped packet. Diagnostic information 230 may include a switch indicator which may identify the switching device where the original packet was dropped. Diagnostic information 230 may include congestion statistics which may include congestion information of the switching device where the original packet was dropped. The intention packet may be generated by the device using software, hardware or a combination thereof.

15

20

Once a switching device generates intention packet 222, the switching device may convey intention packet 222 to the destination device through network 110. Since intention packet 222 may comprise a small packet, it may be assured (or nearly assured) of routing through the subsequent switching devices in network 110 to reach its destination device if congestion was the adverse condition. The switching devices may be configured to insure that they always have the required bandwidth or resources to route

25

packets up to a maximum size. The intention packets may be below this maximum size such that the switching devices may always route them.

5 A device may react to receiving an intention packet based on the information in the intention packet as well as its own characteristics. The information in the intention packet may indicate whether the dropped packet was a request or a response to a request. If the dropped packet was a request, then the device that receives the intention packet, i.e. the receiving device, may send a signal to the requesting device that its request has been dropped. This signal may be a NACK packet conveyed from the receiving device so that
10 the requesting device may resend the dropped packet. The receiving device may also perform an operation that corresponds to the request in anticipation of receiving the resent packet as discussed below in more detail.

15 If the dropped packet was a response to a request, then the requesting device will receive the intention packet. The requesting device may reschedule the request in response to receiving the intention packet and may convey a subsequent packet that includes the request to the receiving device. The requesting device may be configured to optimize the resending of the request by rescheduling the request for a more appropriate time or by merging the request with a subsequent request.

20 In one particular embodiment of Fig. 1, device 100 may comprise a host interface device that is coupled to a computer bus or network (not shown) and device 120 may comprise a disk drive. In this embodiment, the host interface device may perform read and write accesses to the disk drive on behalf of devices in the computer network. The
25 cases of a write request and a read request with the host interface and disk drive will now be discussed to illustrate the use of intention packets in a congested network.

In order to perform a write request, the host interface may convey a packet to the disk drive through network 110. The packet may include a header with information about the write request and a payload that includes one or more data blocks to be written. If each switching device encountered by the packet has sufficient bandwidth to transmit the packet, then the packet can be routed to the disk drive where the disk drive can perform a write operation corresponding to the write request. Depending on the communications protocol employed by the system, the disk drive may convey a packet back to the host interface to indicate that the write request was successful after completing the write operation.

If the bandwidth (or resources) required by the packet is not available at a switching device in network 110, then that switching device may drop the packet and generate an intention packet. The intention packet may include information from the header of the original packet as well as a severed flag and diagnostic information as shown in Fig. 2b. The switching device may then convey the intention packet to the disk drive though network 110. As noted above, the intention packet may be assured (or nearly assured) of avoiding congestion when being routed to the disk drive by virtue of its relatively small size.

When the intention packet is received at the disk drive, the disk drive may convey a NACK to the host interface to indicate that the packet corresponding to the write request has been dropped. The disk drive may record the information regarding the write request such as the location and the size of the write in a list of severed requests, i.e. requests whose packets were dropped in network 110 where the resent requests have not been received by the disk drive. The disk drive may also perform one or more operations corresponding to the write request, such as prefetching old blocks for a parity computation or flushing a write buffer, in anticipation of receiving a resent request from the host interface.

In one embodiment, when the disk drive receives a non-severed request, i.e. a request that is not a resent request, the disk drive may compare the non-severed request to the list of severed requests. If the non-severed request is independent of each of the severed requests, e.g. attempts to read or write a different region of data than each of the severed requests, then the disk drive may complete an operation corresponding to the non-severed request. If the non-severed request is not independent of one or more of the severed requests, then the disk drive may be configured not to complete the operation corresponding to the non-severed request and may convey a NACK to the host interface that sent the non-severed request. For example, each request may have a sequence number. If a non-severed read request for a location arrives at the disk drive and the severed requests list indicates that a severed write request with an earlier sequence number is pending for the same location, then the non-severed read request will be delayed until the severed write request is resent and completed. However, if the non-severed read request is independent from the requests in the severed requests list, then it will be allowed to complete, even if it arrives out of sequence.

When the host interface receives a NACK from the disk drive, either indicating that its request has been dropped or that its request is not independent of another request that has been severed, the host interface may reschedule the request and may convey a new packet corresponding to the request to the disk drive. The host interface may optimize the rescheduled request in a number of ways. For example, the host interface may reschedule the request at a less congested time. In addition, the host interface may merge the request with a subsequent request while ensuring consistency.

25

To perform a read request, a host interface may convey a packet corresponding to the read request to the disk drive through network 110. The disk drive may receive the read request packet and may read the data block or blocks specified by the request. The

disk drive may then convey the data block or blocks in a response packet to the host interface. If each switching device encountered by the response packet has sufficient resources to transmit the packet, then the response packet may be routed to the host interface to complete the read request. If the resources required by the response packet are not available at a switching device in network 110, then that switching device may drop the response packet and generate an intention packet. The intention packet may include information from the header of the response packet as well as a severed flag and diagnostic information, as shown in Fig. 2b. The switching device may then convey the intention packet to the host interface though network 110. Again, the intention packet may be assured (or nearly assured) of avoiding further congestion to the host interface by virtue of its relatively small size.

Note that in some cases, an adverse condition or combination of adverse conditions may be severe enough so that the intention packet cannot be routed or transmitted to its destination or cannot be routed or transmitted to its destination in a reasonable time. Such cases may be handled for the intention packet by conventional mechanisms, such as a time-out mechanism.

When the host interface receives an intention packet, the host interface may reschedule the read request and may convey a new packet corresponding to the read request to the disk drive. The host interface may optimize the rescheduled request in a number of ways. For example, the host interface may reschedule the request after the disk drive has spun around to the location of the request or it may retarget the read request to a less congested mirror copy. In addition, the host interface may merge the request with a subsequent request while ensuring consistency.

Turning now to Fig. 3, a flow chart illustrating a method for implementing a routing scheme employing intention packets in a switching device is shown. Variations

on the method are possible and contemplated. In Fig. 3, a first packet that includes a header and a payload may be received as indicated in block 302. In block 304, a determination may be made if an adverse routing condition is present. For example, a determination may be made as to whether required routing resources (such as bandwidth and buffering space) are available to transmit the payload to a second device. If no adverse condition is detected, then the header may be conveyed to the second device as indicated in block 306 and the payload may be conveyed to the second device as indicated in block 308. It can be noted that the header and payload may be conveyed to the second device at substantially the same time or at separate times depending on a routing model used by a system.

If an adverse condition is detected, such as the required routing resources are not available to transmit the payload to the second device, then the payload may be dropped as indicated in block 310. An intention packet may be generated as indicated in block 312. The intention packet may include all or part of the header of the original packet along with other information as discussed above. The intention packet may be conveyed to the second device as indicated in block 314.

Turning now to Fig. 4, a flow chart illustrating a method for handling packets at a first exemplary device is shown. The first exemplary device may comprise a disk drive or another type of device. Variations on the method are possible and contemplated. In Fig. 4, a packet may be received at a receiving device as indicated in block 402. In block 404, a determination may be made as to whether the packet is an intention packet. In one embodiment, an intention packet may include a flag to indicate that it is an intention packet. If the packet is an intention packet, then intention packet information may be recorded as indicated in block 406. In addition, an operation corresponding to the intention packet may be begun as indicated in block 408, such as prefetching data blocks or flushing a write buffer. A NACK may be conveyed to the sending device as indicated

in block 410. Note that the elements of the method indicated at 408 and 410 do not necessarily occur serially. Element 410 may be performed before 408, after 408 or concurrently with 408.

5 If the packet is not an intention packet, then a determination may be made as to whether an operation corresponding to the packet is independent of an outstanding operation as indicated in block 412. The outstanding operation may correspond to other intention packets that were received at the receiving device prior to the packet. If the operation corresponding to the packet is independent of the outstanding operations, or
10 alternatively, if there are no outstanding operations at the receiving device, then the operation corresponding to the packet may be begun as indicated in block 414. If the operation corresponding to the packet is not independent of the outstanding operations, then a NACK may be conveyed to the sending device as indicated in block 416. In addition, the operation corresponding to the packet may not be performed.

15 Turning now to Fig. 5, a flow chart illustrating a method for handling packets at a second exemplary device is shown. The second exemplary device may comprise a host interface or another type of device. Variations on the method are possible and contemplated. In Fig. 5, a first packet corresponding to a request from a requesting
20 device may be conveyed as indicated in block 502. A second packet corresponding to a response to the request may be received at the requesting device as indicated in block 504. In block 506, a determination may be made as to whether the second packet is an intention packet. In addition, a determination may be made as to whether the second packet is a NACK as indicated in block 508. If the second packet is either an intention
25 packet or a NACK, then the request may be rescheduled as indicated in block 510 and a third packet corresponding to the request may be conveyed from the requesting device as indicated in block 512.

It can be noted that the requesting device may receive an intention packet when the response to its request was dropped in route to the requesting device. The response, for example, may have included data corresponding to a read request from the requesting device. It can be further noted that the requesting device may receive a NACK when its request was not independent of an outstanding request at a receiving device.

Turning now to Fig. 6, a block diagram of one embodiment of a system is shown. Fig. 6 depicts a plurality of application servers 600a through 600(n) (collectively, application servers 600), file system server 610, and disk array 630 coupled to a network 606. Disk array 630 includes controller 632a and a plurality of disk drives 636a through 636(n) (collectively, disk drives 636) coupled to network 634. In addition, controller 632b is coupled to network 606 and disk array 630.

The embodiment of Fig. 6 may be configured to implement the routing scheme discussed above with respect to Figs. 1-5. In particular, controller 632a and controller 632b may each operate as a host interface to network 606 and may perform accesses to disk drives 636 on behalf of application servers 600 and file system server 610. Controller 632a and controller 632b may convey packets to and receive packets from the disk drives 636 through network 634. In response to congestion or other adverse conditions at a switching device in network 634, a packet may be dropped and an intention packet may be generated as discussed above. The intention packet mechanism may also be applied to network 606.

Turning now to Fig. 7, a block diagram of one embodiment of a plurality of nodes connected in by a network is shown. Any number of nodes 710 may be included in the network. In the embodiment of Fig. 7, each node 710 may include a switching device 720. Nodes 710 may also include one or more devices such as a disk drive 730 or local host (e.g. computing) system 740. In other embodiments, each node 710 may include

other types or numbers of devices and the nodes 710 may not be identical. The above-described routing mechanism employing intention packets may be applied to the system of Fig. 7. For example, a device at one node may initiate a request to a device at another node, such as from one host system to another or from a host system to a disk drive. A request may be routed from one node to another through several of the switching devices. If routing problems are encountered, the above-described intention packet routing mechanism may be employed.

Various embodiments may further include receiving or storing instructions and/or data implemented in accordance with the foregoing description upon a carrier medium. Suitable carrier media may include storage media or memory media such as magnetic or optical media, e.g., disk or CD-ROM, as well as transmission media or signals such as electrical, electromagnetic, or digital signals, conveyed via a communication medium such as network 110 and/or a wireless link.

Although the embodiments above have been described in considerable detail, other versions are possible. Numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.